Decoding Task-based Attentional Modulation during Face Categorization

**Yu-Chin Chiu, Michael Esterman, Yuefeng Han,
Heather Rosen, and Steven Yantis**

Abstract

■ Attention is a neurocognitive mechanism that selects task-relevant sensory or mnemonic information to achieve current behavioral goals. Attentional modulation of cortical activity has been observed when attention is directed to specific locations, features, or objects. However, little is known about how high-level categorization task set modulates perceptual representations. In the current study, observers categorized faces by gender (male vs. female) or race (Asian vs. White). Each face was perceptually ambiguous in both dimensions, such that categorization of one dimension demanded selective attention to task-relevant information within the face. We used multivoxel pattern classification to show that task-specific modulations evoke reliably distinct spatial patterns of activity within three face-selective cortical regions (right fusiform face area and bilateral occipital face areas). This result suggests that patterns of activity in these regions reflect not only stimulus-specific (i.e., faces vs. houses) responses but also task-specific (i.e., race vs. gender) attentional modulation. Furthermore, exploratory whole-brain multivoxel pattern classification (using a searchlight procedure) revealed a network of dorsal fronto-parietal regions (left middle frontal gyrus and left inferior and superior parietal lobule) that also exhibit distinct patterns for the two task sets, suggesting that these regions may represent abstract goals during high-level categorization tasks. ■

INTRODUCTION

Attention facilitates processing of task-relevant information (Yantis, 2008; Corbetta & Shulman, 2002; Desimone & Duncan, 1995). Evidence for attentional modulation of cortical activity has been reported in humans using fMRI in multiple perceptual domains. For example, covert visuo-spatial attention modulates activity in the corresponding retinotopic regions of extrastriate cortex (e.g., Kelley, Serences, Giesbrecht, & Yantis, 2008; Yantis et al., 2002). Similar effects of attention upon cortical activity have been observed during attention to visual features (e.g., Liu, Slotnick, Serences, & Yantis, 2003; Saenz, Buracas, & Boynton, 2002) and objects (e.g., Serences, Schwarzbach, Courtney, Golay, & Yantis, 2004; O'Craven, Downing, & Kanwisher, 1999).

These studies have documented attentional modulation of visual properties that are processed or represented in regions that are functionally well-characterized (e.g., retinotopic organization of visual cortex, category selectivity of ventral temporal cortex). However, many common perceptual tasks entail categorization on the basis of complex combinations of visual attributes. The mechanisms of attentional modulation of cortical activity based on high-level perceptual categorization rules are unknown.

Here we devised a face categorization task during which subjects categorized either the gender or the race of ambiguous face morphs. Each stimulus was created by morphing together one Asian male face with one White female face or one Asian female face with one White male face. Therefore, each of them varied in both gender (male vs. female) and race (Asian vs. White; see Figure 1A). We investigated whether these two categorization tasks evoked distinguishable patterns of activity in the cortical face network.

We used face categorization tasks for two reasons. First, race and gender are natural categorization tasks that humans perform frequently, rapidly, and accurately. However, little is known about how people accomplish these socially relevant face categorization tasks. Previous behavioral studies have suggested that different feature components as well as holistic information are critical for the performance of different face categorization tasks (e.g., Smith, Fries, Gosselin, Goebel, & Schyns, 2009; Mangini & Biederman, 2004; Schyns, Bonnar, & Gosselin, 2002). Although allocation of attention to different types of information within the face during the categorization task was suggested as a possible neural mechanism (Schyns et al., 2002), the role of attention in these socially relevant face categorization tasks has yet to be investigated.

Second, the cortical substrate for face perception is well characterized and therefore provides a solid foundation for addressing this question. Three main regions, the inferior occipital face area (OFA; e.g., Allison et al., 1994), the fusiform face area (FFA; e.g., Kanwisher, McDermott, & Chun, 1997; Sergent, Ohta, & MacDonald, 1992), and

Johns Hopkins University

Figure 1. (A) Psychophysical procedure for obtaining the ambiguity threshold. The lower (black) sequence illustrates gender categorization. Using a male-to-female sequence, morph levels 1 through 10 were judged "male," and level 11 was judged "female," at which point the sequence was halted. Later in the session, during a female-to-male sequence (gray) with the same morph set, morph levels 20 through 10 were judged "female," and morph level 9 was judged "male." A male/female ambiguous morph level for this pair of faces was defined as the midpoint of the two judgments or level 10 in this example. The same procedure was applied to Asian/White categorization (gray), and in this example, level 11 was selected as the ambiguous morph for the race categorization task. (B) Face categorization task.
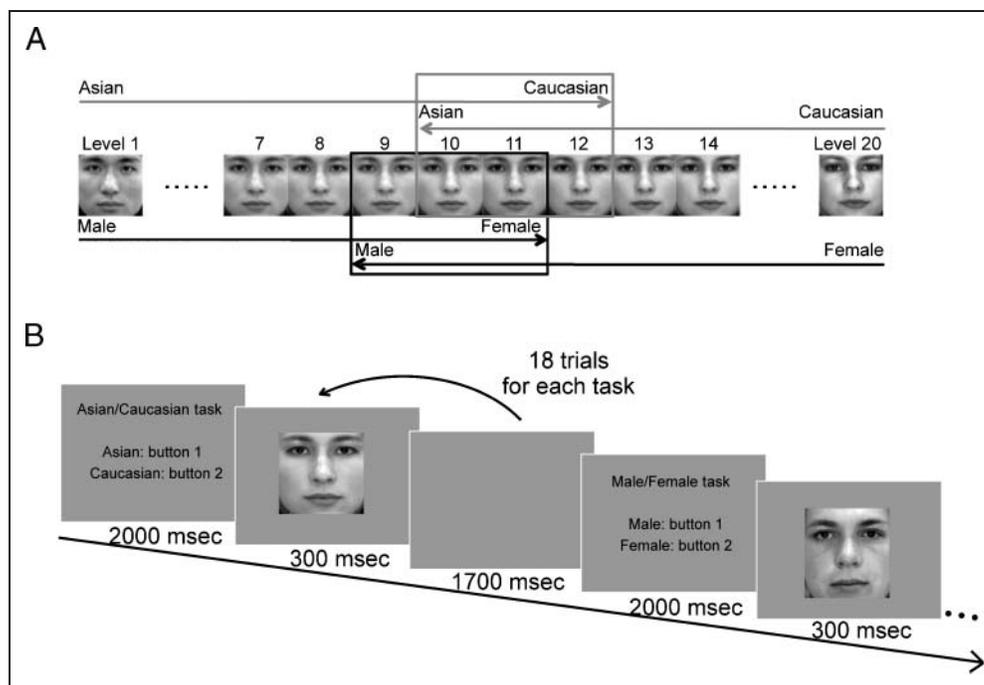
the STS (e.g., Fairhall & Ishai, 2007; Puce, Allison, Bentin, Gore, & McCarthy, 1998), in both hemispheres have been identified as the core face network mediating the visual analysis of human faces and work together as a core system for face perception (Haxby, Hoffman, & Gobbini, 2002).

We employed multivoxel pattern classification (MVPC) in a hypothesis-driven ROI approach to investigate whether cortical face-selective regions reflect not only *what* you see (i.e., responding more to faces than to houses or objects) but also *how* you see it (i.e., responding differently during gender vs. race categorization tasks). In a subsequent exploratory analysis, we used a searchlight procedure (Kriegeskorte, Formisano, Sorger, & Goebel, 2007) to identify additional brain regions exhibiting distinct patterns of activity for the two face categorization tasks.

METHODS

Subjects

Eight neurological healthy adults (all right-handed, four women, age range = 19–33 years, mean = 22 years) gave written informed consent, which was approved by the Johns Hopkins Medicine institutional review board, to participate in this study.

Stimuli and Procedure

Each of six male Asian faces were parametrically morphed (20 levels) with each of six female White faces, and each of six female Asian faces were morphed with each of six male White faces, resulting in a total of 72 pairs of morphs (for an example pair, see Figure 1A). The images were taken from the CalTech database and the AR-face database (Martinez & Benavente, 1998) as well as in-house photography. All stimuli were rendered in grayscale and cropped, leaving only eyes, nose, and mouth, and then morphed using FantaMorph (version 4.0). The faces were presented in the center of the display at a viewing distance of 68 cm and subtended 6° of visual angle in both height and width when viewed in the scanner.

Before the scanning session, subjects completed two sessions on separate days (one categorization task per session) of a preliminary psychophysical task to determine their male–female and Asian–White morph thresholds (i.e., point of subjective ambiguity) for each of the 72 pairs of faces. Subjects were asked to make male/female categorization judgments by pressing one of the two keys on the keyboard. No feedback was provided. Each face was presented until a response was registered. To efficiently obtain subjective thresholds for each pair of stimuli, we used the methods of limits (see Figure 1A). Different morph levels for a pair were presented twice in each session, once in each direction (e.g., once from 100% toward 0% male and once from 0% toward 100% male). The sequence stopped as soon as subjects changed their response (e.g., from male to female). For each pair, the male–female morph threshold was taken as the midpoint of the morph levels that were first categorized as "female" during a male-to-female sequence and as "male" during a female-to-male sequence. Subjects completed the same procedure for the Asian–White morph pairs. The order of tasks in two sessions was counterbalanced across subjects. By the end of the two sessions, two morph levels (one for each task) were determined for

each of the 72 pairs, resulting in a total of 144 ambiguous morphed faces for each subject.

During the scanning session, subjects completed 12, 14, or 16 runs of the face categorization task. Each run consisted of eight blocks of 38 sec each (four blocks of the gender task and four of the race task in alternating order). The initial task in each run was counterbalanced across runs. All 144 stimuli were used in each run, randomly assigned to one of the eight blocks. To ensure that the classification was not based on perceptual differences in the stimuli, we used each face in both tasks across runs. Therefore, the fMRI runs were acquired in pairs: If a face was used in the gender task in the first run of a pair, it was used in the race task in the second run of the pair.

Each of the eight task blocks within a run started with a 2-sec instruction screen indicating which face categorization task to carry out in that block. Following the instruction screen were 18 trials, during each of which a face was presented for 300 msec and a blank screen for 1700 msec (Figure 1B). Subjects held one button in each hand and indicated their categorization decisions by pressing one of the two buttons (response assignments were counterbalanced across subjects).

Independent Functional Localizers

We defined several ROIs to apply MVPC, including six face-selective ROIs (three in each hemisphere; see Figure 2). Each participant completed one or two functional localizer runs lasting 368 sec; these data were used to identify the OFA, the FFA, and the STS in both hemispheres (Figure 2). During each localizer run, subjects alternately viewed 12 blocks (each lasting 30 sec) of intact morphed faces (the 50% morph of 72 pairs) or houses (4 blocks each) and phase-scrambled faces or houses (2 blocks each). Each
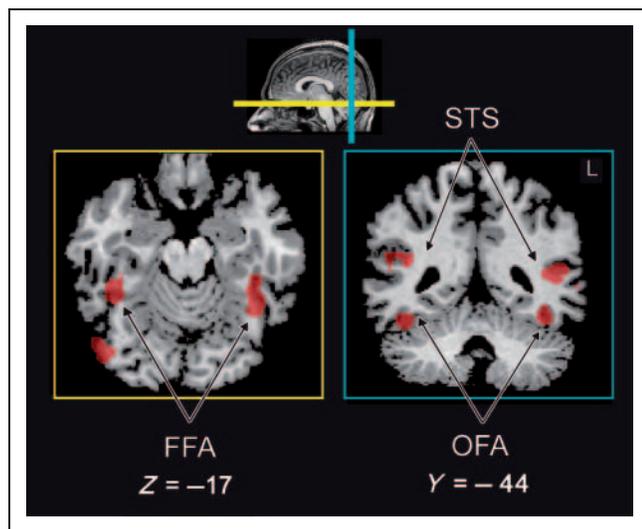

Figure 2. Distributed cortical network for face perception. Data from a representative subject illustrate face-selective ROIs, including OFAs, FFAs, and STS in Talairach space.

image was presented for 300 msec and followed by a 1700-msec blank (Figure 1B). Subjects performed a one-back working memory task and pressed a button when a repetition was detected (three repetitions per block).

fMRI Data Acquisition and Analysis

MRI scanning was carried out on a 3-T Philips Gyroscan scanner. High-resolution anatomical images were acquired with a T1-weighted 200-slice magnetization prepared rapid acquisition gradient-echo sequence with a SENSE (MRI Devices, Inc., Waukesha, WI) eight-channel head coil (repetition time = 8.2 msec, echo time = 3.7 msec, flip angle = 8°, prepulse inversion time delay = 852.5 msec, SENSE factor = 2, scan time = 385 sec), yielding 1-mm isotropic voxels. Whole-brain echo-planar functional images were acquired in 40 transverse slices (repetition time = 2000 msec, echo time = 35 msec, flip angle = 90°, matrix = 64 × 64, with 3 × 3 mm in-plane resolution, slice thickness = 3 mm, SENSE factor = 2).

BrainVoyager QX (Brain Innovation, Maastricht, The Netherlands) and the LIBSVM (Chang & Lin, 2001) for MATLAB were used for analyses. EPI images from each scanning run were slice time and motion corrected and then high-pass filtered (three cycles per run) to remove low-frequency noise in the time series. No spatial smoothing was performed on the images.

Hypothesis-driven ROI MVPC

The localizer data were used to identify face-selective ROIs for each participant using a contrast of faces versus houses and scrambled stimuli. In each ROI, the 50 most selective voxels (see Fox, Iaria, & Barton, 2009) during the preferred stimulation period were included in the subsequent MVPC. The following ROIs were identified, with the number of participants (out of eight total participants) exhibiting significant activation (voxelwise $p < .01$, uncorrected) in that area: right FFA (8), left FFA (7), right OFA (8), left OFA (7), right STS (8), and left STS (7). (One subject had no activation in the vicinity of fusiform gyrus, inferior occipital gyrus, and STS in the left hemisphere with even lowered threshold.) All subsequent ROI analyses were performed using the independent data set from the experimental runs.

The raw time series from each voxel within each ROI was first normalized on a run-by-run basis using a z transform. The mean BOLD signal in each of the 50 voxels in each ROI during a period extending from 12 to 38 sec after the onset of the categorization task instruction screen was taken as one instance for classification. The data were passed through an arctan squashing function to diminish the importance of outliers (Guyon, Weston, Barnhill, & Vapnik, 2002). Thus, four voxel pattern instances were extracted for each task in each run. A leave-one-run-out cross-validation procedure was used to train a linear support vector machine on the basis of all but one run, and the classifier was applied to the data from the left-out run to

discriminate between the gender versus the race categorization tasks. Overall classification accuracy was defined as the mean classification accuracy across all possible permutations of leaving one run out.

Note that the question of interest was whether there exist multivoxel pattern differences in the face network evoked during the gender and race discrimination tasks. Another potentially interesting question would be to assess whether the multivoxel patterns evoked by different *stimuli* (e.g., male vs. female faces or Asian vs. White faces) could be decoded within those ROIs. Although our stimulus set could be usefully applied to this question, the current block design was not optimized for event-related analysis of stimulus differences.

Exploratory Whole-brain MVPC

A 9-mm-cubic searchlight (i.e., 27 voxels) was defined to move through the whole acquired volume (individual native space), centered on each voxel in turn, again using the one-run left out procedure. After obtaining the classification accuracy for each voxel neighborhood in each subject, we applied a Talairach transformation to combine the resulting statistical maps across the eight subjects. The group mean classification accuracy of the searchlight centering on each voxel was then tested against chance (50% accuracy) with a right-tailed t test and corrected for multiple comparisons with a cluster threshold correction (Forman et al., 1995). The final statistical map reported below uses a corrected alpha = .001 with voxelwise nominal p of .003, $t(7) = 3.8$.

RESULTS

Univariate Analyses

We first conducted a univariate ROI-based analysis in the core face network using independent localizer runs (see Methods). We first examined whether the mean BOLD signal in these ROIs differed for the two tasks. To do this, we computed the mean magnitude of the sustained response (i.e., 12–38 sec after the onset of the task instruction screen) during the two task conditions (gender vs. race) across all voxels in each ROI and performed a paired t test between them. None of the ROIs exhibited significant mean differences for the two tasks (all $ps > .14$). This is not surprising; these face-selective areas were, on average, about equally engaged during the two face categorization tasks.

Multivariate Pattern Analyses

A linear support vector machine classifier was trained (separately for each subject) to discriminate multivoxel patterns evoked while subjects carried out the race and gender categorizations, respectively, using data from all the face-selective regions combined for each subject (i.e.,

300 voxels for all subjects except for S8, who contributed 150 voxels, all from the right hemisphere). The mean classification rate was 61% and was significantly greater than chance across subjects, $t(7) = 2.5$, $p < .05$. Second, we trained a classifier for each ROI separately. Figure 3 shows the mean pattern classification rate of the group in each face-selective ROI. Chance classification performance was 50%. Right OFA, right FFA, and left OFA exhibited classification performances that were significantly greater than chance ($ps < .05$), with mean classification accuracies of 61%, 58%, and 57%, respectively.

As another check that the results were not due to chance, we used a randomization procedure in which runs were randomly labeled as gender task or race task before classification. In each region, classification accuracy during cross validation was near 50% (i.e., 48.9–50.8%), as expected. When these values were used as an empirical measure of chance, the same regions (i.e., bilateral OFA and right FFA) exhibited significant classification accuracy.

Although there were no significant overall differences in mean signal across the two tasks in each of these ROIs (see above univariate analyses), to ensure the classification was not driven by mean differences (see Esterman, Chiu, Tamber-Rosenau, & Yantis, 2009), we additionally conducted the same analysis with removal of mean differences. We continued to observe significant classification performance in right OFA, right FFA, and left OFA (63%, 57%, and 57%, respectively, $ps < .05$) but not in other ROIs. The classification rate for all the face-selective regions combined was also significantly better than chance (i.e., 64%, $p < .05$). This mean-centering procedure ensured that a nonspecific difference (e.g., in efforts) was not driving pattern classification; instead, specific multivoxel patterns of activity were reliably different for the two task sets.

We then used an exploratory MVPC searchlight to examine classification on the basis of patterns from $3 \times 3 \times 3 = 27$-voxel clusters centered at each voxel across the whole
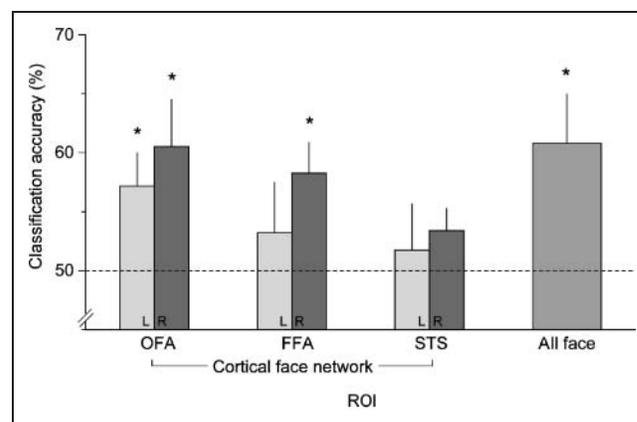

Figure 3. Classification performance in the cortical face-selective ROIs. Mean classification accuracy (%) for each ROI and for all face-selective ROIs combined. Chance is 50%. OFA = occipital face area; FFA = fusiform face area.

brain to identify additional regions that contain distinct signals for the two face categorization tasks. This procedure revealed a set of frontal and parietal regions in the left hemisphere, including middle frontal gyrus, superior and inferior parietal lobule, and insula as well as several clusters in the middle and inferior temporal (IT) cortex (see Table 1). Several of the identified regions in the IT cortex were in and near the fusiform gyrus in both hemispheres as shown in Figure 4, although these were not identical to the independently localized face-selective regions (i.e., FFA, shown in green in Figure 4, and OFA).

Behavioral Analyses

To ensure that the results were not driven by differences in mental effort or personal preference for one category or response versus the other, we analyzed the response times as well as the relative proportion of each response type. Mean response time for the gender and race categorization tasks was 913 and 939 msec, respectively ($p > .13$, data from six subjects; behavioral data from two subjects were lost because of technical issues). There were also no response biases in either task. All eight subjects categorized faces as female as often as male during the gender task (51% vs. 49%, respectively, $p > .8$). The same pattern was observed in the race task (46% vs. 54% for Asian, White, respectively, $p > .1$). These results were as expected because the stimuli used during the scanning sessions were selected to be subjectively ambiguous in both tasks (see
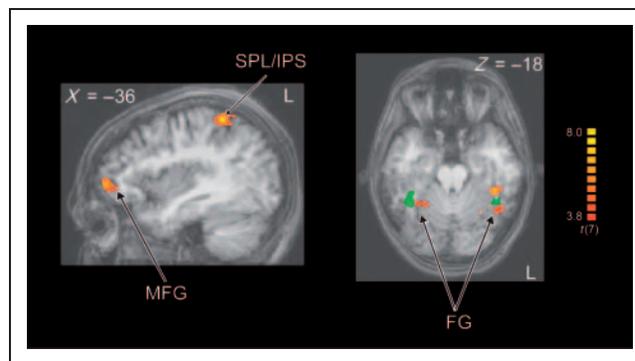

Figure 4. Task-specific signals revealed by the MVPC searchlight. Exploratory whole-brain MVPC revealed regions exhibiting classification that was significantly better than chance for the group (yellow/orange). These included a subset of the fronto-parietal attentional control network, including middle frontal gyrus (MFG), and superior parietal lobule (SPL)/intraparietal sulcus (IPS) as well as several clusters in the IT cortex. Some ventral temporal regions in fusiform gyrus in both hemispheres are adjacent to the group locus of the functionally defined FFA (shown in green). These ROIs are projected onto an averaged anatomical brain in Talairach space.

Methods). In addition, there was no consistent activity that was greater for one task versus the other (see above), further suggesting that the tasks were well matched for difficulty and processing demands.

DISCUSSION

The data reported here reveal modulation of the cortical face network evoked by high-level categorization task set. Critically, because the stimuli were identical during the two tasks, classification could not be based on intrinsic sensory differences in the physical stimuli. Under these conditions, among all the face-selective brain regions, we found three (right FFA and bilateral OFA) that exhibited distinct, task-specific multivoxel patterns evoked by the two task sets. Our results suggest that these regions represent the diagnostic features or combinations of features that are critical for these categorizations. Furthermore, our findings echo studies of prosopagnosic patients or TMS-induced neurodisruption of face-selective areas (for a review, see Rossion, 2008). For example, Minnebusch, Suchan, Köster, and Daum (2009) found that the failure to process faces in subjects suffering from developmental prosopagnosia was linked to the lack of activation in bilateral OFA and FFA. Pitcher, Walsh, Yovel, and Duchaine (2007) found that repetitive TMS of right OFA disrupted accurate discrimination of face parts; however, no effect of repetitive TMS was observed in left OFA. Because the categorization tasks used here relied on high-level face perception rather than on low-level judgments such as size or luminance, our finding of OFA and FFA is in line with previous literature.

This conclusion is also consistent with previous studies, also suggesting that different diagnostic sensory information is critical for different face categorization tasks. Smith

Table 1. Regions Containing Task-specific Signals Revealed by the MVPC Searchlight

| Side | Region | No. Voxels | Peak Voxel | | |
			X	Y	Z
Right	Middle temporal gyrus	432	57	−31	−5
Right	Postcentral gyrus	583	42	−28	34
Right	Insula	598	30	11	1
Right	Fusiform gyrus	455	24	−46	−17
	Cingulate gyrus/ paracentral lobule	920	0	−22	46
	Cuneus	392	0	−82	25
Left	Superior parietal lobule	386	−24	−46	46
Left	Inferior parietal lobule	1134	−36	−43	49
Left	Middle frontal gyrus	466	−36	47	−2
Left	Inferior temporal gyrus	1199	−51	−55	−14
Left	Fusiform gyrus	734	−24	−68	−8
Left	Fusiform gyrus	414	−39	−37	−20
Left	Inferior occipital gyrus/ FFA	619	−33	−76	−11

et al. (2009) and Schyns et al. (2002) used a psychophysical subsampling technique to demonstrate that gender categorization relied more on the eyes and mouth whereas identification relied on almost the whole face. Independently, Mangini and Biederman (2004), using classification images, concluded that different aspects of the face stimuli were critical for different categorization tasks. It is likely that the differences between race and gender categorization depends on both individual face components (e.g., eyes, nose, and mouth) and configural information (e.g., eye separation, eye–nose distance, etc.), but not on other socially relevant information (e.g., gaze), which are thought to be processed in STS (e.g., Fairhall & Ishai, 2007). Furthermore, Sigala and Logothetis (2002) used single neuron recording in monkeys to measure the neural representations of task-specific diagnostic information in the IT cortex following categorization learning (of line drawings of faces). After learning, neurons in IT became tuned to diagnostic aspects of the face (e.g., eye height) to correctly categorize face stimuli.

It is possible that subjects used a purely component-based strategy for discrimination. For example, one could fixate on the eyes during one categorization task and fixate on the mouths during the other categorization task. If this strategy were employed, our MVPC results could be partially driven by low-level visual properties that differ at fixation (e.g., two elliptical contours vs. one). However, this is unlikely for two reasons. First, the searchlight analysis failed to classify patterns in early visual cortex, where distinct foveal stimulation would most likely lead to distinct patterns of responses. Second, regions that did contain distinct multivoxel patterns (i.e., FFA, OFA) are regions known to have large receptive fields, which should be relatively insensitive to small changes in the retinal position of the stimulus. It is possible, however, that regions in intraparietal sulcus (IPS) that support reliable MVPC could be associated with different overt (or covert) states of attention to different facial features because IPS has been associated with different locations of covert attention and saccades (e.g., Schluppeck, Glimcher, & Heeger, 2005).

Although the ROI-based MVPC results reflect task-specific modulation of face-selective cortex (FFA and OFA), the exploratory searchlight approach revealed several fronto-parietal regions (see Figure 4) that contain task-specific patterns of activity. These signals may represent sources of top–down control during task set maintenance, in contrast to the effects of control in target regions like FFA and OFA. The fronto-parietal findings echo many studies suggesting that both pFC and parietal cortex play a role in maintaining behavioral goals, intentions, and abstract task rules (e.g., Badre, 2008; Koechlin & Summerfield, 2007; Miller, 2000). Converging evidence for this idea can be found in several recent studies, including single neuron recording in monkeys (e.g., Asaad, Rainer, & Miller, 2000), human fMRI with univariate analysis (e.g., Bengtsson, Haynes, Sakai, Buckley, & Passingham, 2009; Chiu & Yantis, 2009; Braver, Reynolds, & Donaldson, 2003; Sohn, Ursu, Anderson, Stenger, &

Carter, 2000), and MVPC decoding studies (e.g., Bode & Haynes, 2009; Esterman et al., 2009; Haynes et al., 2007).

The searchlight analysis revealed several regions in IT cortex that partially overlap and neighbor the localizer-defined face-selective regions (see Figure 4). Several reasons for inexact correspondence in these two results are possible. First, the precise anatomical location of face-selective cortex is variable across subjects, and the Talairach transformation is not optimal for combining data across subjects in this case; the searchlight analysis was conducted separately in each subject and then combined. Second, and more importantly, voxels that are revealed by the searchlight reflect visual analysis relevant to categorizations with ambiguous stimuli (i.e., the main gender/race task); in contrast, the face localizer task reflects basic level categorization and stimulus matching (see also Nestor, Vettel, & Tarr, 2008).

In summary, we used MVPC to successfully decode specific face categorization tasks in a subset of the cortical face network, revealing task-specific attentional modulation of face representations. Within the core face network, right FFA and bilateral OFA contained discriminable task-dependent signals. We also observed distinct task-specific signals in the left dorsal fronto-parietal network (i.e., middle frontal gyrus, IPS, superior parietal lobule), which may play a role in abstract goal maintenance. We postulate that these regions may be sources of the task-specific control signals that evoke the distinct patterns observed in the ventral face-selective cortical regions. These results provide further evidence for the importance of FFA and OFA in face processing and expand our knowledge of how top–down attention can flexibly bias information processing to meet task goals.

REFERENCES

Allison, T., Ginter, H., McCarthy, G., Nobre, A. C., Puce, A., Luby, M., et al. (1994). Face recognition in human extrastriate cortex. *Journal of Neurophysiology, 71,* 821–825.

Asaad, W. F., Rainer, G., & Miller, E. K. (2000). Task-specific neural activity in the primate prefrontal cortex. *Journal of Neurophysiology, 84,* 451–459.

Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends in Cognitive Sciences, 12,* 193–200.

Bengtsson, S. L., Haynes, J. D., Sakai, K., Buckley, M. J., & Passingham, R. E. (2009). The representation of abstract task rules in the human prefrontal cortex. *Cerebral Cortex, 19,* 1929–1936.

Bode, S., & Haynes, J. D. (2009). Decoding sequential stages of task preparation in the human brain. *Neuroimage, 45,* 606–613.

Braver, T., Reynolds, J., & Donaldson, D. (2003). Neural mechanisms of transient and sustained cognitive control during task switching. *Neuron, 39,* 713–726.

Chang, C.-C., & Lin, C.-J. (2001). LIBSVM: A library for support vector machines.

Chiu, Y.-C., & Yantis, S. (2009). A domain-independent source of cognitive control for task sets: Shifting spatial attention and switching categorization rules. *Journal of Neuroscience, 29,* 3930–3938.

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3,* 201–215.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18,* 193–222.

Esterman, M., Chiu, Y.-C., Tamber-Rosenau, B. J., & Yantis, S. (2009). Decoding cognitive control in human parietal cortex. *Proceedings of the National Academy of Sciences, U.S.A., 106,* 17974–17979.

Fairhall, S. L., & Ishai, A. (2007). Effective connectivity within the distributed cortical network for face perception. *Cerebral Cortex, 17,* 2400–2406.

Forman, S. D., Cohen, J. D., Fitzgerald, M., Eddy, W. F., Mintun, M. A., & Noll, D. C. (1995). Improved assessment of significant activation in functional magnetic resonance imaging (fMRI): Use of a cluster-size threshold. *Magnetic Resonance in Medicine, 33,* 636–647.

Fox, C. J., Iaria, G., & Barton, J. J. (2009). Defining the face processing network: Optimization of the functional localizer in fMRI. *Human Brain Mapping, 30,* 1637–1651.

Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. *Machine Learning, 46,* 389–422.

Haxby, J. V., Hoffman, E. A., & Gobbini, M. I. (2002). Human neural systems for face recognition and social communication. *Biological Psychiatry, 51,* 59–67.

Haynes, J. D., Sakai, K., Rees, G., Gilbert, S., Frith, C., & Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Current Biology, 17,* 323–328.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience, 17,* 4302–4311.

Kelley, T., Serences, J. T., Giesbrecht, G., & Yantis, S. (2008). Cortical mechanisms for shifting and holding visuospatial attention. *Cerebral Cortex, 18,* 114–125.

Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences, 11,* 229–235.

Kriegeskorte, N., Formisano, E., Sorger, B., & Goebel, R. (2007). Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proceedings of the National Academy of Sciences, U.S.A., 104,* 20600–20605.

Liu, T., Slotnick, S., Serences, J., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex, 13,* 1334–1343.

Mangini, M. C., & Biederman, I. (2004). Making the ineffable explicit: Estimating the information employed for face classification. *Cognitive Science, 28,* 209–226.

Martinez, A. M., & Benavente, R. (1998). *The AR face database.* (CVC Tech. Rep. No. 24).

Miller, E. K. (2000). The prefrontal cortex and cognitive control. *Nature Reviews Neuroscience, 1,* 59–65.

Minnebusch, D. A., Suchan, B., Köster, O., & Daum, I. (2009). A bilateral occipitotemporal network mediates face perception. *Behavioural Brain Research, 198,* 179–185.

Nestor, A., Vettel, J. M., & Tarr, M. J. (2008). Task-specific codes for face recognition: How they shape the neural representation of features for detection and individuation. *PLoS ONE, 3,* e3978.

O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature, 401,* 584–587.

Pitcher, D., Walsh, V., Yovel, G., & Duchaine, B. (2007). TMS evidence for the involvement of the right occipital face area in early face processing. *Current Biology, 17,* 1568–1573.

Puce, A., Allison, T., Bentin, S., Gore, J. C., & McCarthy, G. (1998). Temporal cortex activation in humans viewing eye and mouth movements. *Journal of Neuroscience, 18,* 2188–2199.

Rossion, B. (2008). Constraining the cortical face network by neuroimaging studies of acquired prosopagnosia. *Neuroimage, 40,* 423–426.

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience, 5,* 631–632.

Schluppeck, D., Glimcher, P., & Heeger, D. J. (2005). Topographic organization for delayed saccades in human posterior parietal cortex. *Journal of Neurophysiology, 94,* 1372–1384.

Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science, 13,* 402–409.

Serences, J. T., Schwarzbach, J., Courtney, S. M., Golay, X., & Yantis, S. (2004). Control of object-based attention in human cortex. *Cerebral Cortex, 14,* 1346–1357.

Sergent, J., Ohta, S., & MacDonald, B. (1992). Functional neuroanatomy of face and object processing. A positron emission tomography study. *Brain, 115,* 15–36.

Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature, 415,* 318–320.

Smith, M. L., Fries, P., Gosselin, F., Goebel, R., & Schyns, P. G. (2009). Inverse, mapping the neuronal substrates of face categorizations. *Cerebral Cortex, 19,* 2428–2438.

Sohn, M. H., Ursu, S., Anderson, J. R., Stenger, V. A., & Carter, C. S. (2000). The role of prefrontal cortex and posterior parietal cortex in task switching. *Proceedings of National Academy of Sciences, U.S.A., 97,* 13448–13453.

Yantis, S. (2008). Neural basis of selective attention: Cortical sources and targets of attentional modulation. *Current Directions in Psychological Science, 17,* 86–90.

Yantis, S., Schwarzbach, J., Serences, J., Carlson, R., Steinmetz, M. A., Pekar, J. J., et al. (2002). Transient neural activity in human parietal cortex during spatial attention shifts. *Nature Neuroscience, 5,* 995–1002.